

(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 016 987 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
05.07.2000 Bulletin 2000/27

(51) Int. Cl.⁷: G06F 17/30

(21) Application number: 99204514.6

(22) Date of filing: 23.12.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 29.12.1998 US 222088

(71) Applicant:
SUN MICROSYSTEMS, INC.
Palo Alto, California 94043 (US)

(72) Inventor:
Thulasiraman, Jeyaraman
San Jose, CA 95117 (US)

(74) Representative:
Shortt, Peter Bernard et al
TOMKINS & CO.,
5 Dartmouth Road
Dublin 6 (IE)

(54) Propagating updates efficiently in hierarchically structured data

(57) One embodiment of the present invention provides a system that efficiently propagates changes in hierarchically organized data to remotely cached copies (118) of the data. The system operates by receiving (302) an access to the data at a client. In response to this access, the system determines (304) if the client (106) contains a copy of the data. If so, the system sends a request (308) to a server (104) for an update (122) to the copy. The server (104) receives the request and determines differences (310) between the current version of the data (116) at the server and an older copy of the data (118) at the client, which the server has stored locally. These differences are used to construct (312) an update (122) for the copy of the data, which may include node insertion and node deletion operations for hierarchically organized nodes in the data. Next, the update is sent (314) to the client (106) where it is applied to the copy of the data to produce an updated copy of the data. Finally, the original access is allowed (318) to proceed on the updated copy of the data. According to one aspect of the present invention, the act of determining differences (310), and the act of using the differences to construct (312) the update both take place during a single pass through the data. According to another aspect of the present invention, the update (122) for the copy of the data may include node copy, node move, node collapse and node splitting operations.




FIG. 1

Description

SUMMARY

BACKGROUND

[0001] The present invention relates to distributed computing systems and databases. More particularly, the present invention relates to a method and an apparatus that facilitates detecting changes in hierarchically structured data and producing corresponding updates for remote copies of the hierarchically structured data.

[0002] The advent of the Internet has led to the development of web browsers that allow a user to navigate through inter-linked pages of textual data and graphical images distributed across geographically distributed web servers. Unfortunately, as the Internet becomes increasingly popular, the Internet often experiences so much use that accesses from web browsers to web servers often slow to a crawl.

[0003] In order to alleviate this problem, a copy of a portion of a web document from a web server (document server) can be cached on a client computer system, or alternatively, on an intermediate proxy server, so that an access to the portion of the document does not have to travel all the way back to the document server. Instead, the access can be serviced from a cached copy of the portion of the document located on the local computer system or on the proxy server.

[0004] However, if the data on the document server is frequently updated, these updates must propagate to the cached copies on proxy servers and client computer systems. Such updates are presently propagated by simply sending a new copy of the data to the proxy servers and client computer systems. However, this technique is often inefficient because most of the data in the new copy is typically the same as the data in the cached copy. In this case, it would be more efficient to simply send changes to the data instead of sending a complete copy of the data.

[0005] This is particularly true when the changes to the data involve simple manipulations in hierarchically structured data. Hierarchically structured data typically includes a collection of nodes containing data in a number of forms including textual data, database records, graphical data, and audio data. These nodes are typically inter-linked by pointers (or some other type of linkage) into a hierarchical structure, which has nodes that are subordinate to other nodes, such as a tree -- although other types of linkages are possible.

[0006] Manipulations of hierarchically structured data may take the form of operations on nodes, such as node insertions, node deletions or node movements. Although such operations can be succinctly stated and easily performed, there presently exists no mechanism to transmit such operations to update copies of the hierarchically structured data. Instead, existing systems first apply the operations to the data, and then transmit the data across the network to update copies of the data on local machines and proxy servers.

[0007] One embodiment of the present invention provides a system that efficiently propagates changes in hierarchically organized data to remotely cached copies of the data. The system operates by receiving an access to the data at a client. In response to this access, the system determines if the client contains a copy of the data. If so, the system sends a request to a server for an update to the copy. The server receives the request and determines differences between the current version of the data at the server and an older copy of the data at the client, which the server has stored locally. These differences are used to construct an update for the copy of the data, which may include node insertion and node deletion operations for hierarchically organized nodes in the data. Next, the update is sent to the client where it is applied to the copy of the data to produce an updated copy of the data. Finally, the original access is allowed to proceed on the updated copy of the data. According to one aspect of the present invention, the act of determining differences, and the act of using the differences to construct the update both take place during a single pass through the data. According to another aspect of the present invention, the update for the copy of the data may include node copy, node move, node collapse and node splitting operations.

BRIEF DESCRIPTION OF THE FIGURES

[0008]

FIG. 1 illustrates a computer system including a web browser and a web server in accordance with an embodiment of the present invention.

FIG. 2 illustrates a computer system including a server that automatically updates local copies of documents in accordance with another embodiment of the present invention.

FIG. 3 is a flow chart illustrating how a client requests an update from a server in accordance with an embodiment of the present invention.

FIG. 4 is a flow chart illustrating how a server automatically updates local copies of documents in accordance with an embodiment of the present invention.

FIG. 5 is a flow chart illustrating how the system creates updates for a new copy of hierarchically structured data in accordance with an embodiment of the present invention.

FIGs. 6A-6I illustrate the steps involved in creating updates to transform a document tree T1 into a document tree T2.

DETAILED DESCRIPTION

[0009] The following description is presented to enable any person skilled in the art to make and use the

invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

[0010] The data structures and code described in this detailed description are typically stored on a computer readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital video discs), and computer instruction signals embodied in a carrier wave. For example, the carrier wave may carry information across a communications network, such as the Internet.

Computer System

[0011] FIG. 1 illustrates a computer system including a web browser and a web server in accordance with an embodiment of the present invention. In the illustrated embodiment, network 102 couples together server 104 and client 106. Network 102 generally refers to any type of wire or wireless link between computers, including, but not limited to, a local area network, a wide area network, or a combination of networks. In one embodiment of the present invention, network 102 includes the Internet. Server 104 may be any node coupled to network 102 that includes a mechanism for servicing requests from a client for computational or data storage resources. Client 106 may be any node coupled to network 102 that includes a mechanism for requesting computational or data storage resources from server 104.

[0012] Server 104 contains web server 112, which stores data for at least one web site in the form of inter-linked pages of textual and graphical information. Web server 112 additionally includes a mechanism to create updates for remotely cached copies of data from web server 112.

[0013] Web server 112 stores textual and graphical information related to various websites in document database 116. Document database 116 may exist in a number of locations and in a number of forms. In one embodiment of the present invention, database 116 resides within the same computer system as server 104. In another embodiment, document database resides at a remote location, and is accessed by server 104 through network 102. Note that portions of document database 116 may reside in volatile or non-volatile semiconductor memory. Alternatively, portions of document database 116 may reside within rotating storage

devices containing magnetic, optical or magneto-optical storage media.

[0014] Client 106 includes web browser 114, which allows a user 110 viewing display 108 to navigate through various websites coupled to network 102. Web browser 114 stores cached copies 118 of portions of website documents in local storage on client 106.

[0015] During operation the system illustrated in FIG. 1 operates generally as follows. In communicating with web browser 114, user 110 generates an access to a document in web server 112. In processing the access, web browser 114 first examines cached copies 118 to determine if the access is directed to a portion of a web document that is already cached within client 106. If so, client 106 makes an update request 120, which is transferred across network 102 to server 104. In response to the request, server 104 generates an update 122, which is transferred to web browser 114. Update 122 is then applied to the cached copies 118 in order to update cached copies 118. Finally, the access is allowed to proceed on the cached copies 118.

[0016] Note that although the example illustrated in FIG. 1 deals with web documents for use with web browsers and web servers, in general the present invention can be applied to any type of data. This may include data stored in a hierarchical database. This may also include data related to a directory service that supports a hierarchical name space.

[0017] Also, server 104 and web server 112 may actually be a proxy server that stores data in transit between a web server and web browser 114. In this case, the invention operates on communications between the proxy server and web browser 114.

[0018] In a variation on the embodiment illustrated in FIG. 1, client 106 is a "thin client" with limited memory space for storing cached copies of documents 118. In this variation, when client 106 requests a document, only a subset of the document that client 106 is actually viewing sent from server 104 to client 106. This subset is adaptively updated as client 106 navigates through the document.

[0019] In another variation on the above embodiment, documents from document database 116 are tree-structured. In this variation, documents or portions of documents that are sent from server 104 to client 106 are first validated to ensure that they specify a proper tree structure before they are sent to client 106. This eliminates the need for client 106 to validate the data. (Validation is typically performed by parsing the data, constructing a tree from the data, and validating that the tree is properly structured.) Reducing this work on the client side can be particularly useful for thin clients, which may lack computing resources for performing such validation operations.

[0020] FIG. 2 illustrates a computer system including a server that automatically updates local copies of documents in accordance with another embodiment of the present invention. In the embodiment illustrated in

FIG. 2, network 202 couples together server 204 with workstation 206, personal computer 208, network computer 210 and personal organizer 212. Network 202 generally refers to any type of wire or wireless link between computers, including, but not limited to, a local area network, a wide area network, or a combination of networks. In one embodiment of the present invention, network 202 includes the Internet. Server 204 may be any node coupled to network 202 that includes a mechanism for servicing requests from a client for computational or data storage resources. Server 204 communicates with a number of clients, including workstation 206, personal computer 208, network computer 210 and personal organizer 212. In general, a client may include any node coupled to network 202 that contains a mechanism for requesting computational or data storage resources from server 204. Note that network computer 210 and personal organizer 212 are both "thin clients," because they have rely on servers, such as server 204 for data storage and computational resources. Personal organizer 212 refers to any of a class of portable personal organizers containing computational and memory resources. For example, personal organizer 212 might be a PALMPILOT™ distributed by the 3COM Corporation of Sunnyvale, California. (PalmPilot is a trademark of the 3COM Corporation).

[0021] In the illustrated embodiment, workstation 206, personal computer 208, network computer 210 and personal organizer 212 contain cached documents 207, 209, 211 and 213, respectively. Cached documents 207, 209, 211 and 213 contain locally cached portions of documents from server 204.

[0022] Server 204 is coupled to document database 214, which includes documents to be distributed to clients 206, 208, 210 and 212. Document database 214 may exist in a number of locations and in a number of forms. In one embodiment of the present invention, document database 214 resides within the same computer system as server 204. In another embodiment, document database resides at a remote location that is accessed by server 204 across network 202. Portions of document database 214 may reside in volatile or non-volatile semiconductor memory. Alternatively, portions of document database 214 may reside within rotating storage devices containing magnetic, optical or magneto-optical storage media.

[0023] Server 204 includes publishing code 205, which includes computer code that disseminates information across network 202 to workstation 206, personal computer 208, network computer 210 and personal organizer 212. Publishing code 205 includes a mechanism that automatically creates updates for locally cached copies of documents from document database 214 stored in clients 206, 208, 210 and 212.

[0024] During operation, the system illustrated in FIG. 2 operates generally as follows. Publishing code 205 periodically receives new content 230, and uses

new content 230 to update documents within document database 214. Publishing code also periodically constructs updates for remotely cached copies of documents from document database 214, and sends these updates to clients, such as workstation 206, personal computer 208, network computer 210 and personal organizer 212. Note that these updates do not simply contain new versions of cached documents, but rather specify changes to cached documents.

Updating Process

[0025] FIG. 3 is a flow chart illustrating how a client requests an update from a server in accordance with an embodiment of the present invention. This flow chart describes the operation of the invention with reference to the embodiment illustrated in FIG. 1. First, the system receives a request access the data (step 302). In FIG. 1, this corresponds to user 110 requesting access to a web page or a portion of a web page through web browser 114 on client 106. Next, the system determines if client 106 contains a copy of the data (step 304). This corresponds to web browser 114 looking in cached copies 118 for the requested data. If the data is not present on client 106, the system simply sends a copy of the requested data from server 104 to client 106 (step 306), and this copy is stored in cached documents 118.

[0026] If a copy of the data is present on client 106, client 106 sends an update request 120 to server 104 requesting an update to the copy (step 308). In one embodiment of the present invention, update request 120 includes a time stamp indicating how long ago the previous update to cached documents 118 was created. In response to update request 120, server 104 determines differences between the copy of the data on client 106, and the data from document database 116 (step 310). These differences are used to construct an update 122, which specifies operations to update the copy of the data on client 106 (step 312). Note that if client 106 sends a timestamp along with the request in step 308, the timestamp can be used to determine the differences between the data on server 104 and the cached copy of the data on client 106. In another embodiment of the present invention, server 104 saves update 122, so that server 104 can send update 122 to other clients. In yet another embodiment, server 104 keeps track of changes to the data from document database 116 as the changes occur; these changes are aggregated into update 122. This eliminates the need to actually find differences between the data from document database 116 and the cached copy of the data on client 106.

[0027] Also note that the operations specified by update 122 may include manipulations of nodes with in the data. For example, if the data is hierarchically organized as nodes in a tree structure, the update may specify tree node manipulation operations, such as move, swap, copy, insert and delete operations for leaf nodes. The data may also specify sub-tree move, copy, swap,

insert, delete operations, as well as internal node splitting and internal node collapsing operations. Transmitting such node manipulation operations, instead of transmitting the data that results after the node manipulation operations have been applied to the data, can greatly reduce the amount of data that must be transmitted to update a copy of the data on client 106.

[0028] The update may additionally include a Multipurpose Internet Mail Extensions (MIME) content type specifying that the update contains updating operations for hierarchically organized data. This informs a client receiving update 122 that update 122 contains update information, and not regular data. The MIME content type may specify that update 122 contains updating information that has been validated by server 104 so that client 106 does not have to validate update 122.

[0029] In one embodiment of the present invention, the steps of determining the differences (step 310) and constructing the update 122 (step 312) take place concurrently during a single pass through the data. This technique has performance advantages over performing these steps separately in two passes through the data.

[0030] Next, update 122 is sent from server 104 to client 106 (step 314), and client 106 applies update 122 to the copy of the data (step 316). In one embodiment of the present invention, the copy of the data is stored in semiconductor memory within client 106, and hence applying update 122 to the copy of the data involves fast memory operations, instead of slower disk access operations.

[0031] Finally, the original access to the data (from step 302) is allowed to proceed, so that user 110 can view the data on display 108. The above process is repeated for successive accesses to the copy of the data on client.

[0032] Note that although the illustrated embodiment of the present invention operates in the context of a web browser and a web server, the present invention can be applied in any context where updates to data have to be propagated to copies of the data. For example, the present invention can be applied to distributed database systems.

[0033] FIG. 4 is a flow chart illustrating how server 204 (from FIG. 2) automatically updates local copies of documents in accordance with an embodiment of the present invention. This embodiment is an implementation of a "push" model, in which data is pushed from a server 204 to clients 206, 208, 210 and 212 without the clients having to ask for the data. This differs from a "request" model, in which the clients have to explicitly request data before it is sent as is illustrated in FIG. 1.

[0034] The flow chart illustrated in FIG. 4 describes the operation of the invention with reference to the embodiment illustrated in FIG. 2. First, server 204 receives new content 230 (step 402). This new content 230 may take the form of live updates to document database 214, for example in the form of stock pricing infor-

mation. New content 230 is used to update documents or other data objects within document database 214 on server 204 (step 404).

[0035] Next, publishing code 205 within server 204 determines differences between the data in document database 214 and copies of the data on clients (subscribers) 206, 208, 210 and 212, (step 406). These differences are used to construct updates 216, 218, 220 and 222, which specify operations to change copies of the data on clients 206, 208, 210 and 212, respectively (step 408).

[0036] Updates 216, 218, 220 and 222 may specify operations that manipulate nodes within the data. For example, if the data is hierarchically organized as nodes in a tree structure, updates 216, 218, 220 and 222 may specify tree node manipulation operations, such as move, swap, copy, insert and delete operations for leaf nodes. The data may also specify sub-tree move, copy, swap, insert, delete operations, as well as internal node splitting and internal node collapsing operations. Transmitting such node manipulation operations, instead of transmitting the data after node manipulation operations have been applied to it, can greatly reduce the amount of data that must be transmitted to update copies of the data on clients 206, 208, 210 and 212.

[0037] In one embodiment of the present invention, the steps of determining the differences (step 406) and of constructing updates 216, 218, 220 and 222 (step 408) takes place concurrently during a single pass through the data. This can have a significant performance advantage over performing these steps in two separate passes through the data.

[0038] Next, updates 216, 218, 220 and 222 are sent from server 204 to clients 206, 208, 210 and 212 (step 410), respectively. Clients 206, 208, 210 and 212 apply updates 216, 218, 220 and 222 to their local copies of the data 207, 209, 211 and 213, respectively (step 412). In one embodiment of the present invention, these updates are applied to are applied the local copies 207, 209, 211 and 213 "in memory," without requiring disk accesses. This allows to updates to be performed very rapidly.

[0039] The above process is periodically repeated by the system in order to keep copies of the data on clients 206, 208, 210 and 212 at least partially consistent with the data on server 204. This updating process may be repeated at any time interval from, for example, several seconds to many days.

50 Process of Creating Updates

[0040] FIG. 5 is a flow chart illustrating how the system creates updates at the server for a new copy of hierarchically structured data in accordance with an embodiment of the present invention. This embodiment assumes that the data is hierarchically organized as a collection of nodes in a tree structure. This tree structure includes a root node that can have a number of chil-

dren. These children can also have children, and so on, until leaf nodes, which have no children, are reached. Note that the below-described process for creating updates requires only a single pass through the data. During this single pass the system determines differences between old and new trees and creates corresponding updates to convert the old tree into the new tree. This eliminates the need for a separate time-consuming pass through the data to create updates from differences.

[0041] The system starts with an old tree (old_t) and a new tree (new_t). The system first matches leaf nodes of old_t and new_t (step 502). In doing so, the system may look for exact matches or partial matches of the data stored in the leaf nodes. In the case of partial matches, if the quality of a match is determined to be above a preset threshold, the leaf nodes are considered to be "matched." Next, the system generates deletion operations to remove nodes from old_t which are not present in new_t (step 504).

[0042] In the next phase, the system repeats a number of steps (506, 508, 510 and 512) for ascending levels of the tree. First, for a given level, the system generates node insertion operations for nodes that are present in new_t but not in old_t (step 506). Also, if the position of a node in old_t is different from the position of the same node in new_t, the system generates a move operation, to move the node from its position in old_t to its new position in new_t (step 508). Additionally, if a parent node in old_t does not have all of the same children in new_t, the system generates a node split operation for the parent, splitting the parent node into a first parent and a second parent (step 510). The first parent inherits all of the children that are present in new_t, and the second parent inherits the remaining children. If a parent node in old_t has all of the same children and additional children in new_t, the system generates a node collapse operation to bring all the children together in new_t (step 512).

[0043] Additionally, if all of the children of a first parent in old_t move to a second parent in new_t, the system generates a node collapse operation to collapse the first parent into the second parent so that all of the children of the first parent are inherited by the second parent.

[0044] The system repeats the above-listed steps 506, 508, 510 and 512 until the root of the tree is reached. At this point all of the operations that have been generated are assembled together to create an update that transforms old_t into new_t (step 514).

Example

[0045] Let us consider the example tree illustrated in Figure 6A. This tree may represent a document consisting of sections, paragraphs and individual sentences containing parsable character data. Assume that the document grammar also allows documents to contain

non-character data, say numeric data, as is represented by the leaf node identifier 'd'. All nodes in FIG. 6A include a name (tag), a value, and an associated value identifier. Since the leaf nodes actually contain data, value identifiers are assigned to them before the process starts; whereas, for an internal node, a value identifier is assigned during the comparison process based upon the value of identifiers of the internal node's children. Note that in some embodiments of the present invention, the tree data structure as represented in memory may conform to the World Wide Web Consortium document object model (W3CDOM).

[0046] Additionally, in some embodiments of the present invention, the hierarchically organized data includes data that conforms to the Extensible Markup Language (XML) standard. In other embodiments of the present invention, the hierarchically organized data includes data that conforms HyperText Markup Language (HTML) standard, and other markup language standards.

Notational Semantics

[0047] We represent each leaf node by the path from root node to the leaf node containing the position of each node along the path. Hence, the notation for each of the leaf nodes in FIG. 6A is as follows:

D0.Se0.P0.S0 (left-most node)
 D0.Se0.P0.S1
 D0.Se0.P0.S2
 D0.Se0.P0.S3
 D0.Se0.P1.S0
 D0.Se1.N0
 D0.Se2.P0.S0
 D0.Se2.P1.S0
 D0.Se2.P1.S1
 D0.Se2.P2.S0
 D0.Se2.P2.S1 (right-most node)

The above notation is used to locate and represent any node in the tree, whether it be a leaf node or internal node.

[0048] The notational semantics for each of the tree transformation operations is as follows:

- * MOV(D0.Se0.P0.S2, D0.Se2.P1.S0). In FIG. 6A, this operation moves the leaf node with value identifier 'a'. Note that a similar operation can be used to represent a movement of an internal node. In the case of an internal node, the entire sub-tree moves. Thus, the movement of an individual node or a sub-tree can be an inter-parent move or an intra-parent move.
- * SWP(D0.Se0.P0.S2, D0.Se0.P0.S1). This operation is permitted only in the case of nodes that share a common parent (i.e., intra-parent only). The operation swaps the position of the affected nodes,

under the common parent. In the case of internal nodes, entire sub-trees are swapped.

- $CPY(D0.Se0.P0, D0.Se2.P2)$. This operation replicates a node by making an identical copy on the node. In the case of internal nodes, the entire sub-tree is copied.
- $INS(D0.Se0.P0.S0, a', \{data\})$. This operation inserts a node in the tree at the given position and assigns to it a value identifier a' along with the $\{data\}$. In the case of an internal node, $\{data\}$ assigned contains a null value.
- $DEL(D0.Se0.P0)$. This operation deletes a node and all of its children.
- $SPT(D0.Se0.P0, I)$. This operation splits a parent node into a first node and a second node. All of the children of the parent node starting at position I are transferred to the first node. The remaining children are transferred to the second node. The first node gets the same tag type as the original parent node.
- $CLP(D0.Se0.P0, D0.Se0.P1)$. This operation collapses the contents of a first node and a second node. The resulting node gets the same tag type as the first node. The children of the second become the right-most children of the resulting node.
- $UPD(D0.Se0.P0.S2, \{\delta\})$. This operation specifies a change $\{\delta\}$ to the contents of a leaf node. The $\{\delta\}$ itself describes how to apply (or merge) the change.

[0049] The example described below generates a set of operations to transform an old tree $T1$ (FIG. 6A) into a new tree $T2$ (FIG. 6B). Note that in this example the leaf nodes contain actual data, and the internal nodes simply contain tags which organize and describe the data. There are three phases in the process, including: (1) matching the leaf nodes in $T1$ and $T2$; (2) deleting nodes in $T1$ with no match in $T2$; and (3) modifying or moving nodes the remaining nodes to create $T1$.

Phase 1: Matching Leaf Nodes

[0050] The first step is to generate a unique identifier for each of the leaf nodes in $T2$ based on the content of the leaf node. This can be accomplished by using a hash function to generate a unique identifier for each of the leaf nodes. If two leaf nodes have the same content, then the hash function generates the same identifier. If two leaf nodes have the same identifier, it will not cause problems, because the process uses the root node to leaf node path to identify the individual nodes.

[0051] Next, the process assigns value identifiers to leaf nodes of $T1$. For a given leaf node in $T1$, the process uses a hash function to generate a unique identifier, which matches one of the leaf node identifiers in $T2$. If the identifier generated does not match any of the identifiers in $T2$, then process attempts to find a closest matching leaf node in $T2$, based on some matching criteria. For example, the process may use the Longest

Common Sub-sequence (LCS) algorithm ("Data Structures and Algorithms," Aho, Alfred V., Hopcroft, John E. and Ullman, Jeffrey D., Addison-Wesley, 1983, pp. 189-194) to determine a percentage match between the contents of leaf nodes in $T1$ and $T2$. The matching criterion can be flexible. For example, the matching criterion may specify a minimum of 30% commonality in order for the leaf nodes to be matched.

[0052] Allowing matches to be made on an acceptable matching criteria provides a measure of flexibility. In case a given leaf node's content has been only slightly modified in going from $T1$ to $T2$, the system simply matches the node with its modified version in $T2$. The process subsequently makes the leaf nodes consistent through the $UPD(node, \delta)$ operation. However, if the commonality between leaf nodes being matched does not satisfy the matching criterion, the process assigns a unique value identifier to the leaf node in $T1$, which indicates that the leaf node has been deleted.

[0053] In the worst case, the time complexity of finding a match between the leaf nodes will be $O(K^2)$, where K is the number of unique leaf node identifiers in $T1$ and $T2$. In the best case, where the leaf nodes in $T1$ and $T2$ match in a straightforward manner, the complexity will be $2 \cdot K$. However, the number of changes in a document from one version to another is typically fairly small, in which case only a few leaf nodes need to be matched based on the weak matching criteria.

Phase 2: Deletion phase

[0054] After the matching phase is complete, there may be some leaf nodes in $T1$, which are not matched to nodes in $T2$. These unmatched are deleted as follows.

- For unmatched leaf nodes in $T1$ (from left to right), create a delete operation, such as $DEL(D0.Se2.P2.S0)$.
- Reduce the number of delete operations, by replacing them with sub-tree delete operations, if possible. If all children belonging to a parent are to be deleted, the delete operation of each of the children can be replaced by a single delete operation of the parent node. This involves scanning the deletion list, looking for common parents. If $T1$ has K levels, at most $K-1$ scans are needed to identify a common parent for deletion. Notice that while scanning the i th level, unreduced nodes in the $i+1$ level can be ignored, since they cannot be further reduced.
- After the reductions are performed, the final deletion list is repositioned, because deleting a node at position 0 alters the relative positioning of adjacent nodes. Hence, if two delete operations are to be performed on nodes that have a common parent, then the second delete operation needs to be altered to reflect the change in position of the sec-

ond node to be deleted.

[0055] In the instant example, leaf nodes y, t, h and i in FIG. 6A are unmatched. In accordance with the first step, the system creates following delete operations,

```
DEL(D0.Se0.P0.S0),
DEL(D0.Se0.P0.S1),
DEL(D0.Se2.P2.S0),
DEL(D0.Se2.P2.S1).
```

In the second step, scanning left to right (scan level 4), the system notices that all of D0.Se2.P2's children are to be deleted. By reducing the individual delete operations "DEL(D0.Se2.P2.S0)" and "DEL(D0.Se2.P2.S1)" into a single delete operation of the parent "DEL(D0.Se2.P2)" we are left the following delete operations.

```
DEL(D0.Se0.P0.S0),
DEL(D0.Se0.P0.S1),
DEL(D0.Se2.P2).
```

[0056] Continuing with the level 3 scan, the system notices that the only eligible delete operation for reduction is DEL(D0.Se2.P2), since the other delete operations DEL(D0.Se0.P0.S0) and DEL(D0.Se0.P0.S1) are at level 4. Since D0.Se2.P2's parent has other children which do not participate in the delete operation, the reduction ends at scan level 3.

[0057] In the third step, the system checks to see if applying the first delete operation will affect the relative node position of any other delete operation. This involves looking for nodes having the same parent as the node being deleted. If such a node exists, the system adjusts its node position accordingly. Note that the entire deletion list need not be scanned to identify sibling nodes, because the inherent ordering in the deletion list ensures that deletion operations for sibling nodes will be close together in the deletion list.

[0058] Continuing with the example, the system notices that applying the delete operation DEL(D0.Se0.P0.y0) will affect the relative positioning of sibling node D0.Se0.P0.t1. So, the system adjusts the position of its sibling (See FIG. 6C). Hence, the final deletion list becomes,

```
DEL(D0.Se0.P0.S0),
DEL(D0.Se0.P0.S0),
DEL(D0.Se2.P2).
```

Phase 3: Modification Phase

[0059] The modification phase brings together the children of internal nodes, in a bottom-up fashion. This involves scanning all the nodes from the bottom-most level (furthest from the root), and scanning each level until level zero is reached. Note that the identity of each

internal node is established by the collective identity of its children. For example, if a parent node's children are identified as 'a' and 'b' respectively, then the identity of the parent is 'ab.'

[0060] Also, if a parent node is left with no children as a result of a move operation, the parent node is deleted. Furthermore, in the special case where there is a skewed tree or sub-tree of nodes having just one child, i.e., a->b->c->d, when node 'd' is deleted, node 'c' is also be deleted. This action is repeated until node 'a' is deleted as well. Instead of generating an individual delete operation for each one of the nodes, the chain of delete operations is reduced to a single delete operation of the grandest common parent of all nodes being deleted.

[0061] Pseudo-code for one embodiment of the modification phase appears below. For each level_i in T2 (leaf to the root) {

1. TO_BE_COMPLETED_LIST = list of all the node value identifiers at level_i in T2.

2. If the node in the TO_BE_COMPLETED_LIST is the root node, find the matching node 't' in T1'. If 't' happens is a root node, break from the loop. Else, partition T1' into two nodes, such that the sub-tree rooted at 't' is moved away from T1', and becomes another tree (T1''). Next, delete the source partition (T1') by deleting its grandest common parent (the root). T1'' and T2 are now identical.

} (end of for loop)

3. Pick one of the nodes 'k' from TO_BE_COMPLETED_LIST, typically the left-most node. SIBLING_LIST = siblings of 'k', including 'k'. Note that we use the term 'node' in place of a node identifier, for convenience.

4. If none of the nodes in the SIBLING_LIST have a matching node in T1', create a parent node 'p' in T1', having the same tag type as the one in T2 (i.e. same as the parent of the nodes in the sibling list in T2). Insert all of the nodes in the sibling list into the newly created parent node in T1'. Next, move the newly created node (along with its children) to be the child of any internal node, preferably, one of the parent nodes at level_i-2, if such a level exists.

5. Let S be the subset of nodes in the SIBLING_LIST that have a match in T1'. Find a parent node 'p' in T1', which has the most siblings in the SIBLING_LIST.

• Move the rest of the matched nodes in S, to be the children of 'p'. If any subset of nodes being moved have a common parent 'q', and if 'q' has no other children, then collapse 'q' into 'p'. Else, individual nodes are moved by separate move

operations.

- The unmatched nodes in the SIBLING_LIST are inserted into 'p'.
- Order the children of 'p' through swap operations. At the end of the swaps, all the children of 'p' which do not happen to be the children of its peer, if any, are gathered in the right-most corner. If there are such children, then a node split operation is performed, so that 'p' has exactly the same children as its peer. The newly created node (sub-tree) is at the same level as 'p' and has the same parent as 'p'. Also, the tag type of 'p' is changed to be the same as its peer in T2, if it is different.

6. Assign a node identity to 'p'. which is the collective identity of its latest children. Similarly, assign an identical identity to the peer node of 'p' in T2.

7. TO_BE_COMPLETED_LIST = TO_BE_COMPLETED_LIST-SIBLING_LIST.

8. If TO_BE_COMPLETED_LIST is not equal to NULL, then return to step 2, else continue.

[0062] Note that the above node movement operations cause changes in the relative positioning of sibling nodes. Hence, the node operations generated by the process should take into account the positional changes caused by node movements.

[0063] The system now applies the modification algorithm on T1' from FIG 6C.

Level 3 scan

[0064] Applying steps 1 and 2, TO_BE_COMPLETED_LIST = {g, c, f, e, b, a, z} and SIBLING_LIST = {g, c}. The system locates the children 'g' and 'c' in T1', and chooses D0.Se2.P1 to be the parent. Applying step 4, the system notices that nodes D0.Se2.P1 and D0.Se0.P1 need to be collapsed. This brings together all the nodes in the SIBLING_LIST under a common parent (See FIG. 6D).

CLP(D0.Se2.P1, D0.Se0.P1)

[0065] Next, the system uses swap operations to re-order the nodes (see FIG. 6E),

SWP(D0.Se2.P1.S1, D0.Se2.P1.S0)
SWP(D0.Se2.P1.S2, D0.Se2.P1.S1)

[0066] Next, a split operation is performed to move away children which do not truly belong to the parent (see FIG. 6F)

SPT(D0.Se2.P1, 2)

[0067] Applying step 5, the system generates an identity for D0.Se2.P1 and its peer in T2 (see FIG. 6G). Though T2 is not shown, it is assumed that the identity has been assigned

[0068] Applying step 6, the system determines that TO_BE_COMPLETED_LIST = {f, e, b, a, z}. Since TO_BE_COMPLETED_LIST is not empty, the system returns to step 2. SIBLING_LIST = {f}. Step 3 and 4 do not produce any changes. Step 5 assigns an identity to D0.Se2.P2. Step 6 removes 'f' from TO_BE_COMPLETED_LIST. Repeating the same, the system eliminates 'd', and 'e' from TO_BE_COMPLETED_LIST.

[0069] At this point, TO_BE_COMPLETED_LIST = {b, a, z} and SIBLING_LIST = {b, a, z}. Step 4 selects node D0.Se0.P0 as a matching node. At this point, node 'z' in the SIBLING_LIST is unmatched in T2. Hence, the system inserts node 'z'. Next, the system applies swap operations to order the children of D0.Se0.P0. Now, TO_BE_COMPLETED_LIST is NULL (see FIG. 6H).

INS(D0.Se0.P0.S0, z, {data})
SWP(D0.Se0.P0.S2, D0.Se0.P0.S0)

Level 2 scan

[0070] Applying steps 1 and 2 the system determines TO_BE_COMPLETED_LIST = {gc, f, e, baz} and SIBLING_LIST = {gc, f}. Applying step 4, the system chooses D0.Se2 as the parent. The system next applies swap operations to order the children, and then split the parent D0.Se2 to move away children that do not belong to D0.Se2. Applying step 5, the system generates identities for D0.Se2 and its peer in T2 (see FIG. 6I).

SWP(D0.Se2.P0, D0.Se2.P1)
SWP(D0.Se2.P1, D0.Se2.P2)
SPT(D0.Se2, 2)

[0071] Now, TO_BE_COMPLETED_LIST = {e, baz} and SIBLING_LIST = {e, baz}. Applying step 4, the system chooses D0.Se0 as the parent. Since, P(e) is the only child, the system collapses D0.Se0 and D0.Se3, and then re-orders the children through swap operations. Applying step 5, the system generates identities for D0.Se0 and its peer in T2 (see FIG. 6J).

CLP(D0.Se0, D0.Se3)
SWP(D0.Se0.P0, D0.Se0.P1)

Level 1 scan

[0072] Applying steps 1 and 2 the system determines TO_BE_COMPLETED_LIST = {ebaz, d, gc} and SIBLING_LIST = {ebaz, d, gc}. Step 4 selects D0 as the parent, and applies the swap operations to re-order its children, which produces T2 (see FIG 6K).

SWP(D0.Se0, D0.Se1)

[0073] Hence, the final set of transformations to transform T1 to T2 is:

DEL(D0.Se0.P0.S0),
 DEL(D0.Se0.P0.S0),
 DEL(D0.Se2.P2),
 CLP(D0.Se2.P1, D0.Se0.P1),
 SWP(D0.Se2.P1.S1, D0.Se2.P1.S0),
 SWP(D0.Se2.P1.S2, D0.Se2.P1.S1),
 SPT(D0.Se2.P1, 2),
 INS(D0.Se0.P0.S0, z, {data}),
 SWP(D0.Se0.P0.S2, D0.Se0.P0.S0),
 SWP(D0.Se2.P0, D0.Se2.P1),
 SWP(D0.Se2.P1, D0.Se2.P2),
 SPT(D0.Se2, 2),
 CLP(D0.Se0, D0.Se3),
 SWP(D0.Se0.P0, D0.Se0.P1), and
 SWP(D0.Se0, D0.Se1).

[0074] Additionally, if partial matches of leaf nodes were made, the leaf nodes need to be updated using UPD operations.

[0075] The above process requires all nodes in T2 be visited and matched with corresponding nodes in T1 once. The complexity of matching the internal nodes is $O(n1+n2)$, where $n1$ and $n2$ are the internal node counts of T1 and T2, respectively. Note that nodes can be matched by hashing node value identifiers.

[0076] Node movements and modifications also add to the overhead. If we consider a cost-based analysis, the cost of a transformation operation on a node 'Y' is a function of the number of children of 'Y'. Thus, the net cost of all transformations will be a function of the total number of nodes involved directly or indirectly in the transformation.

[0077] Since there are no cycles in the transforming operations, the overhead contributed by the node movements is bounded by $O(LK)$, where L is the number of levels in the tree, and K is a the number of leaf nodes. However, typically the number of nodes involved in the movements is very small and does not involve all the nodes in a tree.

[0078] Hence, the worst case time complexity of the algorithm is a summation of the cost of matching leaf nodes $O(K^2)$, the cost of matching internal nodes $O(n1+n2)$, and overhead contributed by node movements $O(LK)$. In an average case analysis, where the number of changes to a document are less than, for example, 20%, the time complexity is a summation of, the cost of matching leaf nodes $O(K)$, the cost of matching internal nodes $O(n1+n2)$, and overhead contributed by node movements $O(K)$.

Optimizations

[0079] There exist a number of additional optimiza-

tions that can be applied to the above process.

- While trying to find a parent 'p' in T1' which has the most children in the SIBLING_LIST, if there is tie, choose a parent with the same tag-type as the one in T2.
- While re-ordering nodes within the same parent (intra-node movement) through swap operations, if the node being moved out is not in the SIBLING_LIST, it can be directly moved to be the right-most child.
- While re-ordering nodes within the same parent (intra-node movement) through swap operations, if the node being moved out is in the SIBLING_LIST, try to position the node being moved out through another swap operation.

[0080] The foregoing descriptions of embodiments of the invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the invention to the forms disclosed. Many modifications and variations will be apparent to practitioners skilled in the art. Accordingly, the above disclosure is not intended to limit the invention. The scope of the invention is defined by the appended claims.

Claims

1. A method for propagating changes in data that is hierarchically organized to a copy (118) of the data, comprising:

receiving (302) , at a client (106), a request to access to the data;
 determining (304) if the client (106) contains the copy (118) of the data;
 if the client (106) contains the copy (118), sending (308) a request (120) to a server (104) for an update to the copy of the data;
 after the request (120) is sent to the server (104), receiving (314) from the server (104) an update (122) for the copy of the data, wherein the update may include node insertion and node deletion operations for hierarchically organized nodes in the data;
 applying (316) the update (120) to the copy of the data to produce an updated copy of the data; and
 allowing (318) the requested access to the data to proceed.

2. The method of claim 1, wherein the update (122) additionally includes at least one from the group of node move, node collapse, node split and node update operations.

3. The method of claim 1 or claim 2, wherein the act of

- sending (308) the request to the server (104) includes sending an indicator of when the copy of the data was last updated.
4. The method of any one of claims 1 to 3, wherein the update includes a Multipurpose Internet Mail Extensions (MIME) content type specifying that the update (122) contains updating operations for hierarchically organized data.
 5. The method of any one of claims 1 to 4, wherein the act of applying (316) the update (122) to the copy (118) of the data takes place in semiconductor memory, whereby the update is able to proceed rapidly in the absence of time-consuming I/O operations.
 6. A method for propagating changes in data that is hierarchically organized to a copy (118) of the data, comprising:
 - receiving (308), at a server (104), a request for an update for the copy of the data (118) located on a client (106);
 - in response to the request, determining differences (310) between the data and the copy of the data;
 - using the differences to construct (312) the update (122) for the copy of the data, wherein the update may include node insertion and node deletion operations for hierarchically organized nodes; and
 - sending (314) the update (122) from the server (104) to the client (106).
 7. The method of claim 6, wherein the acts of determining the differences (310) and constructing (312) the update (122) take place during a single pass through the data.
 8. The method of any one of claims 6 to 7, wherein:
 - the act of sending (308) the request to the server (104) includes sending an indicator of when the copy of the data was last updated; and
 - the act of determining differences (310) includes examining the indicator to determine how the data on the server has changed since the copy of the data was last updated.
 9. The method of any one of claims 6 to 8, wherein the update (122) includes a Multipurpose Internet Mail Extensions (MIME) content type specifying that the update contains updating operations for hierarchically organized data.
 10. The method of any one of claims 6 to 9, wherein the update (122) may include one or more of the following:
 - a) a node copy operation that makes an identical copy of a node as well as any subtree of the node that may exist,
 - b) a node move operation that moves a node to another location in a tree of hierarchically organized nodes,
 - c) a node split operation that splits a node into two separate nodes, and divides any children of the node that may exist between the two separate nodes,
 - d) a node collapse operation that collapses two nodes into a single node, which inherits any children of the two nodes that may exist,
 - e) a node deletion operation includes that deleting any nodes that are subordinate to the node,
 - f) a node swap operation that swaps two nodes as well as any subtrees of the nodes that may exist.
 11. The method of any one of claims 6 to 10, wherein the data that is hierarchically organized includes data that conforms to one of the following:
 - a) the HyperText Markup Language (HTML) standard,
 - b) the Extensible Markup Language (XML) standard.
 12. The method of any one of claims 6 to 11, wherein the data that is hierarchically organized includes a hierarchical database.
 13. The method of any one of claims 6 to 12, wherein the data that is hierarchically organized includes a directory service that supports a hierarchical name space.
 14. The method of any one of claims 6 to 13, further comprising receiving (402) changes (230) to the data on the server (204) from an external source.
 15. The method of any one of claims 6 to 14, wherein the copy of the data (118) located on the client (106) includes a copy of a subset of the data, which is preferably adaptively changed as required by the client (106).
 16. The method of any one of claims 6 to 15, wherein the server (104) includes a proxy server for caching data in transit between a server (104) and a client (106).
 17. The method of any one of claims 6 to 16, wherein the update includes data that is validated at the

server.

18. A computer readable storage medium storing instructions that when executed by a computer cause the computer to perform the method of any one of claims 6 to 17. 5
19. An apparatus that propagates changes in data that is hierarchically organized to a copy of the data, comprising: 10
 - a receiving mechanism (112), on a server (104), that receives a request (120) for an update (122) for the copy of the data (118) located on a client (106); 15
 - an update creation mechanism (310, 312), on the server (104), that determines differences between the data and the copy of the data and uses the differences to create the update for the copy of the data, wherein the update may include node insertion and node deletion operations for hierarchically organized nodes in the data; and 20
 - a sending mechanism (314), on the server (104), that sends the update (122) from the server to the client. 25
20. The apparatus of claim 19, wherein the update creation mechanism (310, 312) is configured to determine differences and create the update during a single pass through the data. 30
21. The apparatus of claim 19 or claim 20, wherein the update additionally includes at least one from the group of node move, node collapse, node split and node update operations. 35
22. The apparatus of any one of claims 19 to 21, wherein the update creation mechanism (310, 312) is configured to determine when the copy of the data on the client was last updated, and to determine how the data on the server (104) has changed since the copy of the data (118) was last updated. 40
23. The apparatus of any one of claims 19 to 22, wherein the update includes a Multipurpose Internet Mail Extensions (MIME) content type specifying that the update contains updating operations for hierarchically organized data. 45
24. An apparatus that propagates changes in data that is hierarchically organized to a copy of the data (118), comprising: 50
 - a request generation mechanism (302, 304, 308), on a client, that receives an access to the data (118), determines if the client contains the copy of the data, and if so, sends a request for

an update to a server (104);

an update receiving mechanism (314), on the client, that receives an update for the copy of the data from the server, wherein the update may include node insertion and node deletion operations for hierarchically organized nodes in the data; and

an updating mechanism (316, 318), on the client, that applies the update to the copy of the data on the client to produce an updated copy of the data.

25. A method for propagating changes in data, the data being organized isomorphically to a hierarchy having a plurality of nodes, the method comprising:

receiving (308) at a server (104) a request for an update (122) to a copy of the data (118);
 responsively to the request, determining (310) at the server (104) a set of one or more differences between the copy and a preferred version of the data, the server (104) having access to the copy and the preferred version at least for purposes of determining the set of differences;

using the set of differences to construct (312) an update (122) for the copy, the update (122) being suitable to conform the copy to the preferred version of the data, the update including at least one operation selected from the group of inserting a node in the data hierarchy or deleting a node from the data hierarchy; and making the update thus constructed available for further use.

26. The method of claim 25, wherein the acts of determining the differences (310) and constructing (312) the update take place during a single pass through the data.

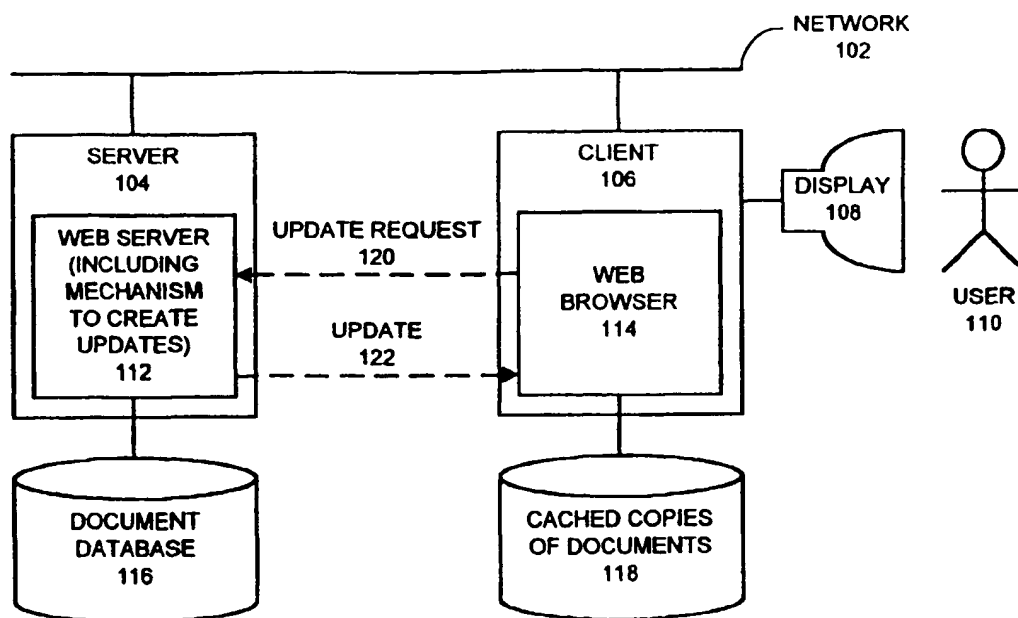


FIG. 1

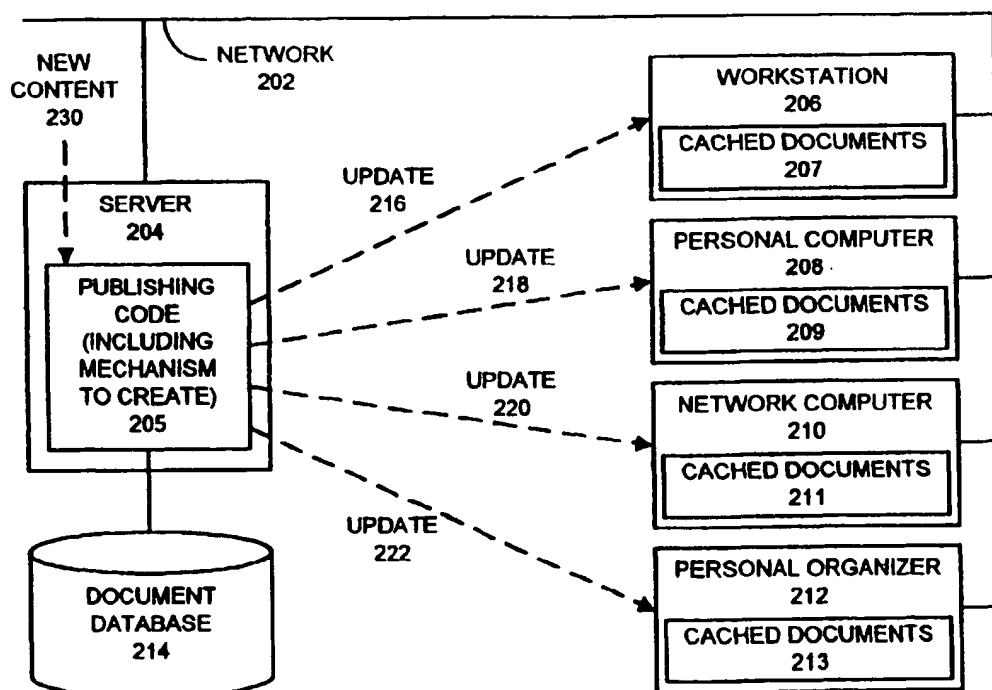


FIG. 2

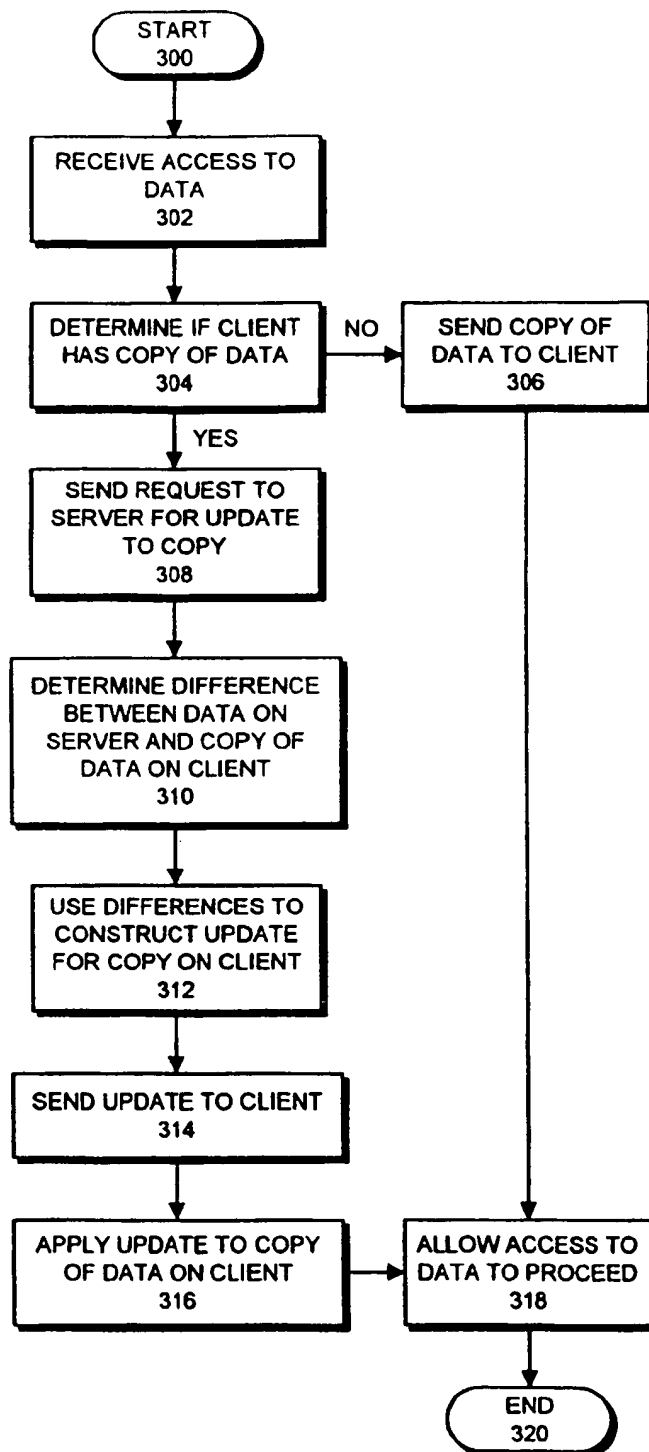


FIG. 3

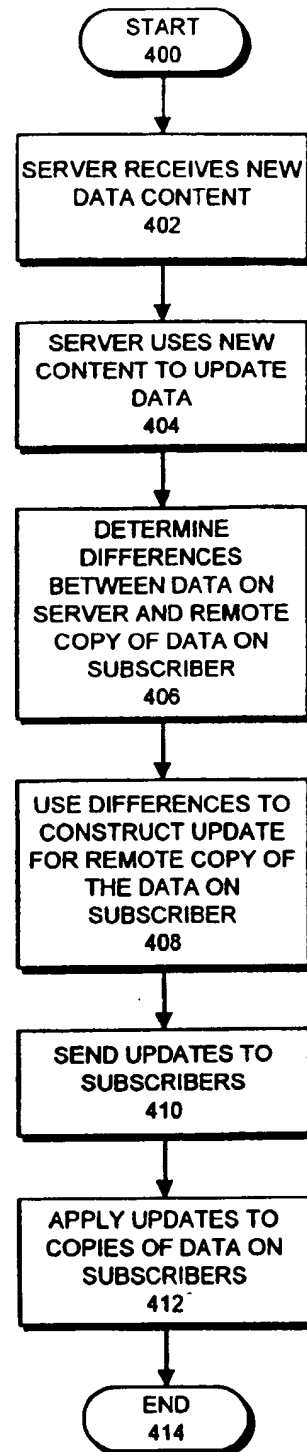


FIG. 4

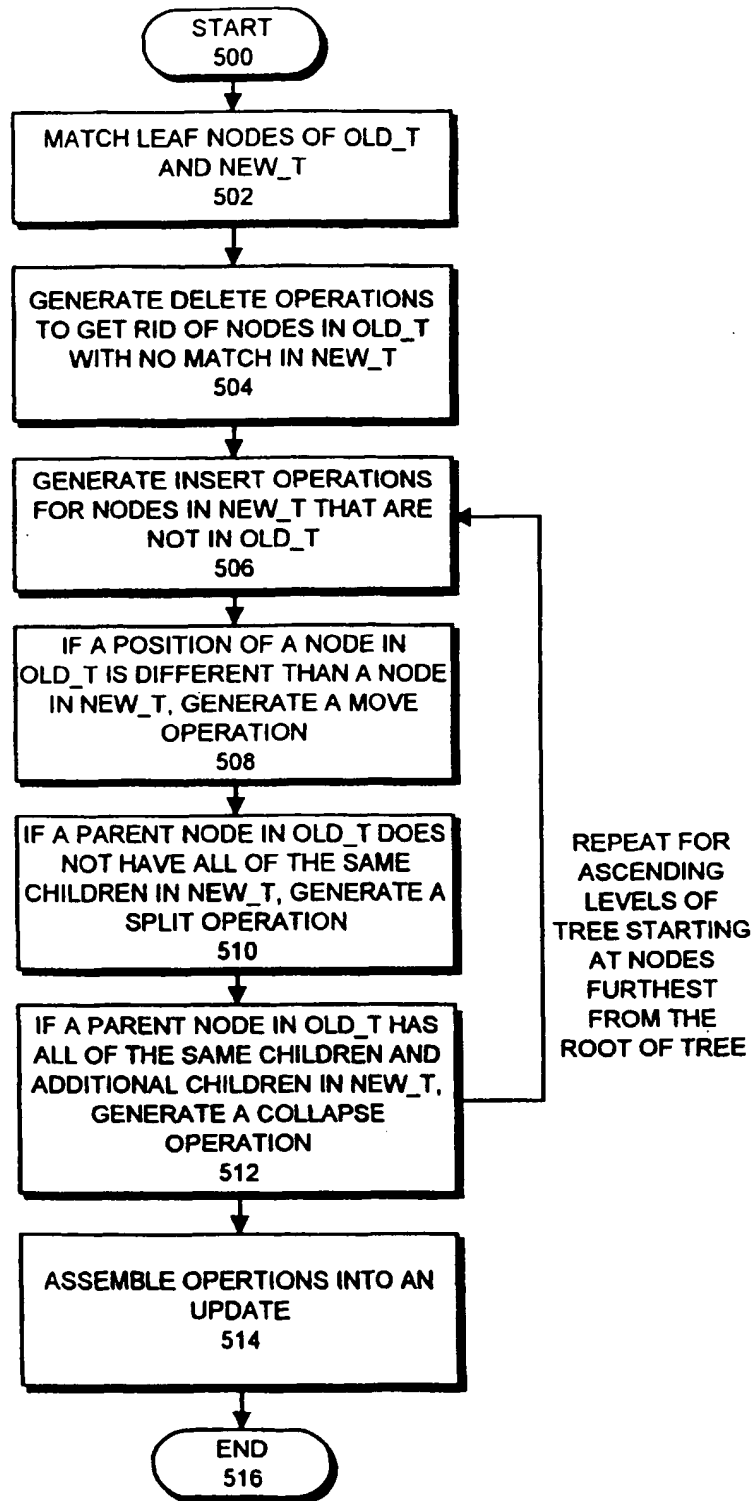


FIG. 5




Figure 6A: Sample Document Tree (T1)




Figure 6B: Modified Document Tree (T2)




Figure 6C: Document Tree (T1') after deletion phase




FIG. 6D




FIG. 6E




FIG. 6F




FIG. 6G




FIG. 6H




FIG. 6I




FIG. 6J





FIG. 6K



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 016 987 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
14.03.2001 Bulletin 2001/11

(51) Int. Cl.⁷: G06F 17/30

(43) Date of publication A2:
05.07.2000 Bulletin 2000/27

(21) Application number: 99204514.6

(22) Date of filing: 23.12.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor:
Thulasiraman, Jeyaraman
San Jose, CA 95117 (US)

(74) Representative:
Hanna, Peter William Derek et al
Tomkins & Co.,
5 Dartmouth Road
Dublin 6 (IE)

(30) Priority: 29.12.1998 US 222088

(71) Applicant:
SUN MICROSYSTEMS, INC.
Palo Alto, California 94043 (US)

(54) Propagating updates efficiently in hierarchically structured data

(57) One embodiment of the present invention provides a system that efficiently propagates changes in hierarchically organized data to remotely cached copies (118) of the data. The system operates by receiving (302) an access to the data at a client. In response to this access, the system determines (304) if the client (106) contains a copy of the data. If so, the system sends a request (308) to a server (104) for an update (122) to the copy. The server (104) receives the request and determines differences (310) between the current version of the data (116) at the server and an older copy of the data (118) at the client, which the server has stored locally. These differences are used to construct (312) an update (122) for the copy of the data, which may include node insertion and node deletion opera-

tions for hierarchically organized nodes in the data. Next, the update is sent (314) to the client (106) where it is applied to the copy of the data to produce an updated copy of the data. Finally, the original access is allowed (318) to proceed on the updated copy of the data. According to one aspect of the present invention, the act of determining differences (310), and the act of using the differences to construct (312) the update both take place during a single pass through the data. According to another aspect of the present invention, the update (122) for the copy of the data may include node copy, node move, node collapse and node splitting operations.




FIG. 1

EP 1 016 987 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 20 4514

DOCUMENTS CONSIDERED TO BE RELEVANT

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
Y	BANGA G ET AL: "Optimistic deltas for WWW latency reduction" PROCEEDINGS OF THE USENIX 1997 ANNUAL TECHNICAL CONFERENCE, PROCEEDINGS OF USENIX 1997 ANNUAL TECHNICAL CONFERENCE, ANAHEIM, CA, USA, 6-10 JAN. 1997, pages 289-303, XP000874445 1997, Berkeley, CA, USA, USENIX Assoc, USA ISBN: 1-880446-84-7 * page 296, left-hand column, line 6 - line 51; figure 6; table 1 * ---	1-3,5-8, 10,11, 14-22, 24-26	G06F17/30
Y	CHANG G J S ET AL: "A graphical environment for change detection in structured documents" PROCEEDINGS TWENTY-FIRST ANNUAL INTERNATIONAL COMPUTER SOFTWARE AND APPLICATIONS CONFERENCE (COMPSAC'97) (CAT. NO.97CB36112), PROCEEDINGS TWENTY-FIRST ANNUAL INTERNATIONAL COMPUTER SOFTWARE AND APPLICATIONS CONFERENCE (COMPSAC'97), WASHINGTON, DC, US, pages 536-541, XP002157517 1997, Los Alamitos, CA, USA, IEEE Comput. Soc, USA ISBN: 0-8186-8105-5 * page 536, right-hand column, line 31 - page 537, left-hand column, line 31 * * page 537, right-hand column, line 31 - page 539, left-hand column, line 20; figures 1-4 * --- -/-	1-3,5-8, 10,11, 14-22, 24-26	<div>TECHNICAL FIELDS SEARCHED (Int.Cl.7)</div> <div>G06F</div>
The present search report has been drawn up for all claims			
Place of search		Date of completion of the search	Examiner
BERLIN		17 January 2001	Polzer, A
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons a : member of the same patent family, corresponding document	
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			

EPO FORM 1503 03 82 (P4/C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 20 4514

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	HOUSEL C ET AL: "WebExpress: a System for Optimizing Web Browsing in a Wireless Environment" PROCEEDINGS OF THE SECOND ANNUAL CONFERENCE ON MOBILE COMPUTING AND NETWORKING, WHITE PLAINS, NY, USA, 11 - 12 November 1996, pages 108-116, XP002157547 * page 110, left-hand column, line 1 - right-hand column, line 23; figure 2 * * page 112, left-hand column, line 1 - page 113, left-hand column, line 8; figures 4-6 *	1,3,6,8, 11, 14-19, 22,24,25	TECHNICAL FIELDS SEARCHED (Int.Cl.7)
A	WILLIAMS S: "HTTP: Delta-Encoding Notes" [Online] 17 January 1997 (1997-01-17) XP002157520 Retrieved from the Internet: <URL: http://ei.cs.vt.edu/~williams/DIFF/prelim.html> [retrieved on 2001-01-16] * section 4.c. *	1,4,6,9, 11,18, 19,23-25	
A	CHAWATHE S S ET AL: "Meaningful change detection in structured data" SIGMOD 1997. ACM SIGMOD INTERNATIONAL CONFERENCE ON MANAGEMENT OF DATA, TUCSON, AZ, USA, 13-15 MAY 1997, vol. 26, no. 2, pages 26-37, XP002157519 SIGMOD Record, June 1997, ACM, USA ISSN: 0163-5808 * page 26, left-hand column, line 1 - page 27, left-hand column, line 9 * * page 27, right-hand column, line 4 - page 28, left-hand column, line 21; figure 1 *	1,2,6,7, 10,11, 18-21, 24-26	
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 17 January 2001	Examiner Polzer, A
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 (03.02) (P04C01)

THIS PAGE BLANK (USPTO)